

Spatially adaptive, Bayesian estimation for probabilistic temperature forecasts

Annette Möller

University of Göttingen

Thordis L. Thorarinsdottir, Alex Lenkoski

Norwegian Computing Center, Oslo

Tilmann Gneiting

Heidelberg Institute for Theoretical Studies and Karlsruhe Institute of Technology

June 16, 2016

Abstract

Uncertainty in the prediction of future weather is commonly assessed through the use of forecast ensembles that employ a numerical weather prediction model in distinct variants. Statistical postprocessing can correct for biases in the numerical model and improves calibration. We propose a Bayesian version of the standard ensemble model output statistics (EMOS) postprocessing method, in which spatially varying bias coefficients are interpreted as realizations of Gaussian Markov random fields. Our Markovian EMOS (MEMOS) technique utilizes the recently developed stochastic partial differential equation (SPDE) and integrated nested Laplace approximation (INLA) methods for computationally efficient inference. The MEMOS approach shows good predictive performance in a comparative study of 24-hour ahead temperature forecasts over Germany based on the 50-member ensemble of the European Centre for Medium-Range Weather Forecasting (ECMWF).

1 Introduction

The two major sources of uncertainty in numerical weather prediction lie in the formulation of the physics based model and in the choice of initial and boundary conditions. These are commonly addressed by using ensemble prediction systems that generate probabilistic forecast ensembles, where the members vary in the details of the numerical model and/or initial and boundary conditions [[Gneiting and Raftery, 2005](#), [Leutbecher and Palmer, 2008](#)].

Statistical postprocessing is employed to correct for biases and dispersion errors in forecast ensembles, based on training data from the past. Specifically, ensemble model output statistics (EMOS) and ensemble Bayesian model averaging (BMA) are widely used postprocessing techniques that yield full predictive distributions from ensemble output; for recent reviews, see, e.g., [Wilks and Hamill \[2007\]](#), [Scheffzik et al. \[2013\]](#), and [Gneiting and Katzfuss \[2014\]](#). Ensemble BMA assigns a kernel density to each bias-corrected ensemble member and combines them into a mixture distribution, using weights that reflect the skill of the individual members. In contrast, EMOS uses a parsimonious distributional regression framework, where a single

parametric predictive distribution is obtained, with the parameters depending on the ensemble members in suitable ways. [Raftery et al. \[2005\]](#) and [Gneiting et al. \[2005\]](#) developed these methods for forecasts of temperature and pressure, with the Gaussian model supplying the underlying kernel density and parametric predictive distribution, respectively. For other weather quantities, alternative distributions are required, as discussed by [Scheffzik et al. \[2013\]](#) and the references therein.

The main objective of statistical postprocessing is to correct for systematic shortcomings in the numerical model. These biases may differ from location to location, due to, e.g., incomplete resolution of the orography or land use characteristics by the model grid. Similarly, the prediction uncertainty may vary over space in ways not represented by the ensemble spread. Consequently, the initial Global EMOS approach [[Raftery et al., 2005](#), [Gneiting et al., 2005](#)], which uses parameters that are spatially constant, entails lesser predictive performance than a Local EMOS approach [[Thorarinsdottir and Gneiting, 2010](#)], which estimates a predictive model at any observation station, based on training data at the station alone. However, the Global approach allows for predictions at just any location, and this is important, as there is an acute need for statistical postprocessing on gridded domains [[Mass et al., 2008](#), [Glahn et al., 2009](#)]. In this light, [Kleiber et al. \[2011a,b\]](#) developed a geostatistical approach that estimates ensemble BMA parameters at each location separately and, subsequently, interpolates them to arbitrary locations by employing a spatial statistical model. [Scheuerer and Büermann \[2014\]](#) and [Scheuerer and König \[2014\]](#) developed a locally adaptive EMOS approach, where information about the short-term local climatology is incorporated and an approach based on intrinsic Gaussian random fields is used to interpolate away from observation locations.

In this paper we propose a Bayesian, locally adaptive implementation of the EMOS method where the regression coefficients for the mean of the predictive distribution vary in space according to Gaussian random fields (GRFs). In order to exploit the most recent available training data, inference for the postprocessing parameters must be repeated in every time step, and thus needs to be computationally efficient. To this end, [Lindgren et al. \[2011\]](#) provide an explicit Markov random field representation of GRFs by solving a certain stochastic partial differential equation (SPDE) on a discretized domain. We utilize this approach in concert with the integrated nested Laplace approximation (INLA) technique [[Rue et al., 2009](#)] to efficiently fit our model in a Bayesian fashion, thereby taking account of estimation uncertainty. Due to the Markovian structure of the GRF, we call our new method Markovian EMOS (MEMOS). While Bayesian inference has been investigated before for BMA [[Vrugt et al., 2008](#), [Di Narzo and Cocchi, 2010](#)], the MEMOS method provides a Bayesian approach to EMOS, and the first Bayesian approach to spatially adaptive postprocessing in general, to our knowledge.

The paper is organized as follows. Section 2 introduces our spatially adaptive, Bayesian estimation approach and its implementation in the SPDE-INLA framework. A case study on forecasts of surface temperature over Germany based on the 50-member ensemble of the European Centre for Medium-Range Weather Forecasts (ECMWF) is presented in Section 3. We end the paper in Section 4 with a discussion of possible extensions of our method.

2 A spatially adaptive, Bayesian approach to inference

2.1 Ensemble model output statistics (EMOS)

The ensemble model output statistics (EMOS) methodology was introduced by [Gneiting et al. \[2005\]](#) as a Gaussian distributional regression technique. Given an ensemble forecast f_1, \dots, f_m for a univariate quantity Y , EMOS employs the ensemble member forecasts as predictors in the

linear model

$$Y = a + b_1 f_1 + \cdots + b_m f_m + \varepsilon,$$

where a, b_1, \dots, b_m are real-valued regression coefficients and $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. For ensembles with exchangeable members, such as the aforementioned ECMWF ensemble, one sets $b_1 = \cdots = b_m$, so that the linear model can be written as

$$Y = a + b\bar{f} + \varepsilon, \tag{1}$$

where $\bar{f} = \frac{1}{m} \sum_{k=1}^m f_k$ is the ensemble mean and, again, $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. In what follows, we focus the discussion on this setting. The parameters a, b , and σ^2 are estimated from forecast and observation data in a rolling training period. [Gneiting et al. \[2005\]](#) compared maximum likelihood estimation to optimizing the continuous ranked probability score (CRPS; see eq. (6) below) and concluded that the latter yields superior predictive performance.

In practice, we need to consider predictions over a domain D that corresponds to the geographic region at hand. To this end, we write $\bar{f}(\mathbf{s})$ and $Y(\mathbf{s})$ to denote the ensemble mean and the observation at the location $\mathbf{s} \in D$, respectively. In the Global EMOS approach the statistical parameters are held constant across space, in that

$$Y(\mathbf{s}) | \bar{f}(\mathbf{s}) \sim \mathcal{N}(a + b\bar{f}(\mathbf{s}), \sigma^2),$$

thereby allowing for large, spatially composited training sets, and yielding estimates with low variance but strong local biases. In the Local EMOS technique, the statistical parameters differ from site to site, so that

$$Y(\mathbf{s}) | \bar{f}(\mathbf{s}) \sim \mathcal{N}(a(\mathbf{s}) + b(\mathbf{s})\bar{f}(\mathbf{s}), \sigma^2(\mathbf{s})),$$

where the statistical parameters are estimated from training data at the location $\mathbf{s} \in D$ only. This yields estimates with small local biases but high variances. As noted, the Global approach allows for predictions at just any location, and geostatistical approaches admit the interpolation of Local EMOS parameters from spatially scattered stations to just any location.

2.2 MEMOS

Our Markovian EMOS (MEMOS) technique interprets the spatially varying bias parameters $a(\mathbf{s})$ and $b(\mathbf{s})$ as realizations of GRFs, using the SPDE-INLA framework for computationally efficient Bayesian inference. This allows the coefficients to adapt to local conditions, while using training data from all observation sites, thereby combining the benefits of both the Global and the Local approaches, and outperforming either.

Specifically, let $\{Y(\mathbf{s}) : \mathbf{s} \in D\}$ denote the future temperature field over the domain D , and let $D_0 \subset D$ denote the combined set of observation and prediction locations. The MEMOS approach utilizes a latent Gaussian Markov random field (GMRF) representation for the bias parameters $\{a(\mathbf{s}) : \mathbf{s} \in D_0\}$ and $\{b(\mathbf{s}) : \mathbf{s} \in D_0\}$ along with Bayesian inference in the SPDE-INLA framework. The result is a joint posterior distribution of the spatially varying bias parameters and the spatially constant variance term σ^2 , given the training data. Conditional on the values of $\bar{f}(\mathbf{s}), a(\mathbf{s}), b(\mathbf{s})$, and σ^2 , the predictive distribution at $\mathbf{s} \in D_0$ is Gaussian,

$$Y(\mathbf{s}) | \bar{f}(\mathbf{s}), a(\mathbf{s}), b(\mathbf{s}), \sigma \sim \mathcal{N}(a(\mathbf{s}) + b(\mathbf{s})\bar{f}(\mathbf{s}), \sigma^2).$$

Integrating over the joint posterior of the postprocessing parameters, we obtain the posterior predictive distribution of the future temperature as a Gaussian variance-mean mixture [Barndorff-Nielsen et al., 1982]. We approximate this posterior by the finite mixture distribution

$$\frac{1}{n} \sum_{i=1}^n \mathcal{N}(a_i(\mathbf{s}) + b_i(\mathbf{s})\bar{f}(\mathbf{s}), \sigma_i^2), \quad (2)$$

where $(a_1(\mathbf{s}), b_1(\mathbf{s}), \sigma_1), \dots, (a_n(\mathbf{s}), b_n(\mathbf{s}), \sigma_n)$ is a sample from the posterior distribution of the random vector $(a(\mathbf{s}), b(\mathbf{s}), \sigma)$. In practice, we represent the predictive distribution by a sample of size $N = mn$ from the finite mixture distribution in (2), namely

$$x_{ij}(\mathbf{s}) = a_i(\mathbf{s}) + b_i(\mathbf{s})\bar{f}(\mathbf{s}) + \sigma_i z_j, \quad i = 1, \dots, n, \quad j = 1, \dots, m, \quad (3)$$

where z_j denotes the standard normal quantile at level $(2j - 1)/(2m)$. The advantages of the specific form of the sample in (3) will become apparent in Section 2.4. In our case study, we use this approach with $m = 50$ and $n = 100$, resulting in a sample of size $N = 5000$.

Before we describe our approach to Bayesian inference using SPDE-INLA, we review some basic facts about GMRFs. As is well known, a stationary GRF $\{X(\mathbf{s}) : \mathbf{s} \in D \subseteq \mathbb{R}^d\}$ with Matérn covariance function [Matérn, 1986, Guttorp and Gneiting, 2008] is a solution to the linear fractional SPDE

$$(\kappa^2 - \Delta)^{\alpha/2} (\tau X(\mathbf{s})) = W(\mathbf{s}), \quad (4)$$

where Δ is the Laplace operator and W is Gaussian white noise with unit variance. Here $\kappa > 0$ is a length scale, $\tau > 0$ is proportional to the marginal standard deviation, and $\alpha = \nu + \frac{d}{2}$ is a smoothness parameter, where $\nu > 0$. When α is an integer the field has the continuous Markov property [see, e.g., Rue and Held, 2005] and admits a GMRF representation. We have tested models with $\alpha = 1$ and $\alpha = 2$, where the latter corresponds to a smoother field, and have found that $\alpha = 1$ yields better predictive performance, so from now on we fix $\alpha = 1$.

Lindgren et al. [2011] proposed an approximate solution of the SPDE (4) by discretizing the spatial domain and, subsequently, solving the equation at the nodes of the grid. As our observations are irregularly spaced, we employ a mesh with triangular grid cells. The algorithm for the mesh creation uses the observation locations as the initial set of nodes, and iteratively adds new nodes until all the triangles in the mesh fulfill a set of regularity conditions, with an equilateral triangle being the most regular. The approximate solution to the SPDE is then the weighted sum

$$X_K(\mathbf{s}) = \sum_{k=1}^K w_k \psi_k(\mathbf{s}) \quad (5)$$

of continuous, piecewise linear basis functions $\psi_k(\mathbf{s})$ with support on the triangles that are attached to the k th node. The random variables w_k are jointly normal with a covariance matrix that depends on the Matérn parameters $\kappa > 0$ and $\tau > 0$.

2.3 Bayesian inference using SPDE-INLA

To perform Bayesian inference we use the SPDE-INLA approach of Lindgren et al. [2011] and Rue et al. [2009] as implemented in the R-INLA package [R Core Team, 2013, Blangiardo et al., 2013, Lindgren and Rue, 2015]. At each estimation-prediction step, the observation locations in the training data are merged with the prediction sites in order to construct the INLA mesh. The model has the hyperparameter vector

$$(\kappa_a, \tau_a, \kappa_b, \tau_b, \sigma),$$

where κ_a and τ_a , and κ_b and τ_b , are the Matérn parameters for the GMRF representation of $\{a(\mathbf{s}) : \mathbf{s} \in D_0\}$, and $\{b(\mathbf{s}) : \mathbf{s} \in D_0\}$, respectively. Our prior assumes that $\log \kappa_a$ and $\log \kappa_b$ are normal with mean -0.082 and variance 1.5 , that $\log \tau_a$ and $\log \tau_b$ are normal with mean -0.878 and variance 1.5 , and that $1/\sigma^2 \sim \text{Gamma}(1, 0.00005)$, with the components being independent. The intercept term $a(\mathbf{s})$ is estimated in two parts; the overall mean level is estimated as a fixed effect, with a random effect describing the local divergence from the mean.

As a result, R-INLA generates samples $\{a_i\}_{i=1}^n$, $\{b_i\}_{i=1}^n$, and $\{\sigma_i\}_{i=1}^n$ from the marginal posterior distributions of $a(\mathbf{s})$, $b(\mathbf{s})$, and σ , respectively, or linear combinations thereof, given the training data. In particular, it generates samples from the marginal posterior distributions of σ and $a(\mathbf{s}) + b(\mathbf{s})\bar{f}(\mathbf{s})$, where \mathbf{s} varies over the prediction sites and \bar{f} is the ensemble mean for the future time point of interest, from which we obtain the approximate predictive distribution (2) and its sample version (3). An important feature of the SPDE-INLA approach is that the discrete representation (5) defines a field at every location $\mathbf{s} \in D$. This is a very useful property for statistical postprocessing directly on the grid on which the underlying numerical model operates, as is frequently required in practice [Mass et al., 2008, Glahn et al., 2009].

2.4 Ensemble copula coupling

It is important to note that the approximate posterior predictive distribution (2) and its sample version (3) apply at each site $\mathbf{s} \in D_0$ independently. Further effort is needed to generate a sample of physically realistic, spatially coherent, postprocessed forecast fields. In the traditional EMOS approach, an attractive option is to fit a geostatistical model to the residual process $\{\varepsilon(\mathbf{s}) : \mathbf{s} \in D\}$ in the basic linear model (1), as proposed by Gel et al. [2004] and recently developed by Feldmann et al. [2015]. Empirical copula based techniques provide attractive, computationally efficient alternatives, including both ensemble copula coupling (ECC) and the Schaake shuffle [Scheffzik et al., 2013, Wilks, 2015].

Here we pursue the latter approach and propose a slight extension of the ECC technique. In a nutshell, ECC restores the rank order structure of the raw ensemble, which results from a sophisticated model of the physics and chemistry of the atmosphere, so the adoption of its dependence structure typically leads to an improvement in the predictive performance. Specifically, let $f_1(\mathbf{s}), \dots, f_m(\mathbf{s})$ denote the raw ensemble forecast, and let $x_1(\mathbf{s}), \dots, x_m(\mathbf{s})$ be a sample from the postprocessed predictive distribution for $Y(\mathbf{s})$, where \mathbf{s} varies over the set D_1 of forecast locations. The sample can be generated in various ways, and in our case study we employ the ECC-Q technique of Scheffzik et al. [2013], which uses equally spaced quantiles of the univariate postprocessed predictive distributions. Finally, let $\pi(\mathbf{s})$ denote the permutation of the integers $1, \dots, m$ that is induced by the order statistics of $f_1(\mathbf{s}), \dots, f_m(\mathbf{s})$, with any ties resolved at random. The ECC ensemble at $\mathbf{s} \in D_1$ is then given by

$$g_1(\mathbf{s}) = x_{(\pi(\mathbf{s})(1))}(\mathbf{s}), \quad \dots, \quad g_m(\mathbf{s}) = x_{(\pi(\mathbf{s})(m))}(\mathbf{s}).$$

While at every location $\mathbf{s} \in D_1$ individually, this is just a reordering, the ECC forecast fields $\{g_j(\mathbf{s}) : \mathbf{s} \in D_1\}$, where $j = 1, \dots, m$, inherit the spatial rank order structure of the raw ensemble. The basic ECC approach thus produces a postprocessed ensembles of the same size m as the raw ensemble, and this is what we do for Global EMOS and Local EMOS. For MEMOS, the Bayesian nature and the specific structure of the sample in (3) allows us to apply the basic ECC-Q procedure to each of the n subsamples $x_{i1}(\mathbf{s}), \dots, x_{im}(\mathbf{s})$ individually. Merging the n respective ECC-Q samples, we obtain a postprocessed ECC ensemble of $N = mn$ spatially coherent forecast fields. Of course, the raw ensemble is invariant under ECC; it already possesses the rank order structure employed in the reordering.

For the purpose of comparison, we assess the predictive performance of the ECC ensembles relative to Independence ensembles, obtained by random permutations of the ensemble members at each location separately.

2.5 Evaluation methods

The goal of probabilistic forecasting is to maximize the sharpness of the predictive distributions subject to calibration [Gneiting et al., 2007]. Here we assess calibration via rank histogram and probability integral transform (PIT) histograms, and we use proper scoring rules as omnibus performance measures.

As our probabilistic forecasts take various forms, it is important to use evaluation tools that allow for meaningful comparison. The raw ensemble forecast is the discrete, empirical distribution associated with its m members. The univariate Local and Global EMOS predictive distributions in (1) are Gaussian, and the univariate MEMOS predictive distribution is the empirical distribution of the sample (3) of size $N = mn$. In the case of multivariate probabilistic forecasts at several locations simultaneously, the Local/Global EMOS and MEMOS predictive distributions are discrete, corresponding to samples of size m and N , respectively.

To assess the calibration of a univariate ensemble of size m , we use the verification rank histogram, which plots the frequency of the rank of the observation, when pooled with the respective ensemble members [see, e.g., Wilks, 2011]. For Gaussian predictive distributions, we use the probability integral transform (PIT) histogram [Dawid, 1984, Gneiting et al., 2007]. The PIT is simply the value of the predictive cumulative distribution function evaluated at the verifying observation. To assess the calibration of the univariate MEMOS ensemble, we compute the verification rank and normalize to the unit interval. In the multivariate case, we employ the multivariate rank histogram described by Gneiting et al. [2008]. In all these cases, calibrated forecasts generate statistically uniform histograms, and deviations from uniformity can be interpreted diagnostically. U-shaped histograms indicate underdispersion, inverse U-shaped histograms suggest overdispersion, and skewed histograms tend to be associated with biases.

Proper scoring rules provide summary measures of predictive performance that may address calibration and sharpness simultaneously [Gneiting and Raftery, 2007]. We take them to be negatively oriented penalties, i.e., the smaller, the better. In the case of univariate predictive distributions, we use the continuous ranked probability score (CRPS). Given the predictive cumulative distribution function, F , and the verifying observation, y , the CRPS is defined as

$$\begin{aligned} \text{CRPS}(F, y) &= \int_{-\infty}^{\infty} (F(x) - \mathbb{1}\{x \geq y\})^2 dx \\ &= \mathbb{E}_F |X - y| - \frac{1}{2} \mathbb{E}_F |X - X'|, \end{aligned} \quad (6)$$

where $\mathbb{1}\{\cdot\}$ denotes the indicator function and X and X' are independent with distribution F [Gneiting and Raftery, 2007]. Closed form expressions are available when F is parametric, such as for Gaussian distributions and finite mixtures thereof [Grimmett et al., 2006]. The representation (6) is valid when F has a finite first moment and implies that the CRPS can be reported in the same unit as the observation. When F is an empirical measure, the expectations become discrete sums, as described by Gneiting et al. [2008], and when F is a deterministic forecast, i.e., a point measure, the CRPS reduces to the familiar absolute error (AE). In computing the AE, we follow Gneiting [2011] and use the Bayes predictor, namely, the median of the respective predictive distribution, as point forecast.

In the multivariate setting, where the predictive distribution is for a quantity in \mathbb{R}^d , we use the proper energy score (ES), which is defined as

$$\text{ES}(F, \mathbf{y}) = \mathbb{E}_F \|\mathbf{X} - \mathbf{y}\| - \frac{1}{2} \mathbb{E}_F \|\mathbf{X} - \mathbf{X}'\|, \quad (7)$$

where $\|\cdot\|$ denotes the Euclidean norm and \mathbf{X} and \mathbf{X}' are independent random vectors with distribution F and finite first moment [Gneiting and Raftery, 2007]. This provides a direct generalization of the CRPS in the form (6). Again, when F is an empirical measure, the expectations become discrete sums.

To compare and rank the various types of forecasts, we find the mean CRPS, AE, or energy score, respectively. In the multivariate setting (Section 3.4) we obtain, for each method, a time series of scores, and for binary key comparisons we apply the Diebold and Mariano [1995] procedure to test the null hypothesis of equal predictive performance in the form of a vanishing expected score differential. In the univariate setting (Section 3.3) we report mean scores both over time and over spatially scattered observation locations. Hering and Genton [2011] proposed a spatial prediction comparison test (SPCT), which is a modification of the Diebold-Mariano test in the time series setting. While this is a very general test, the adaption to our setting is not straightforward, and we leave it to future research. Here, we apply the standard Diebold-Mariano test to the time series of the daily means of the CRPS and the AE, respectively, and furthermore to the time series of the scores at individual sites.

3 Temperature forecasts over Germany

We now present the results of a case study.

3.1 Ensemble forecast and observation data

The data in our case study comprises observations and 24-hour ahead ensemble forecasts of surface (2 meter) temperature over Germany from February 2, 2010 to April 30, 2011. The forecast ensemble is the operational core ensemble of the European Centre for Medium-Range Weather Forecasts (ECMWF), which has $m = 50$ exchangeable members [Molteni et al., 1996, Hemri et al., 2014]. The forecasts are valid at 00 Coordinated Universal Time (UTC), corresponding to 1 am local time in winter, and 2 am local time during the daylight saving period. Observations at 518 stations during the considered time period were obtained and supplied by the German Weather Service (DWD). As the ECMWF forecasts are issued on a grid, at 31 km resolution, they are bilinearly interpolated from the four surrounding grid points to the station locations of interest. The unit used is degrees Celsius.

Panel (a) of Figure 1 shows the spatially scattered locations of the meteorological observation stations in Germany. The station density is highest in urban regions, whereas offshore and outside Germany there are very few stations only. The triangulation is flexible, using triangles of different sizes and adapting locally to the density of the observation locations. To give an example, panel (b) shows the mesh created to estimate the MEMOS model from training data up to October 2, 2010, and used to issue predictions valid October 3, 2010. In creating the mesh, we refrain from extending the data region and require triangles with interior angles of at least 0.1 degrees. The resulting 536 vertices include most, but not all, of the 508 stations in the training set at hand.




Figure 1: (a) Location of the observation stations over Germany. The eleven stations in the Northern coastline example in Section 3.4 are marked in red. (b) Triangulation of the spatial domain for predictions valid October 3, 2010.

3.2 Choice of training period

There is a tradeoff in the choice of an appropriate rolling training period for estimating the MEMOS parameters. Short training periods adapt quickly to seasonally varying biases; longer training periods provide more data and reduce the statistical variability in the estimation. In the extant literature, rolling training periods of length between 20 and 40 days have been common practice. Here we investigate lengths from 15 to 50 days in five day increments over a common test period ranging from March 24, 2010 to April 30, 2011. For MEMOS, both the mean CRPS and the mean AE are minimal under a training period of 25 days. For Local EMOS, longer training periods are slightly favored, but result in negligible improvement only. In this light, we use a rolling training period of length 25 days for all methods. For Local EMOS, in some instances observations are available at a subset of the 25 most recent calendar days only, and then we use the most recent 25 days for which data are available. Furthermore, we fix March 24, 2010 to April 30, 2011 as evaluation period.

3.3 Results at individual stations

We begin with a comparison of the Raw ECMWF to the postprocessed Global EMOS, Local EMOS, and MEMOS forecasts at individual stations. Table 1 shows the mean CRPS and mean AE, averaged over the evaluation period and all available sites, for a total of 205,572 forecast cases. The postprocessed Global EMOS forecast improves substantially on the Raw ECMWF ensemble, as it corrects for biases and dispersion errors. While Global EMOS estimates only a single set of parameters for all stations, Local EMOS estimates a separate set of parameters at each considered station, and Local EMOS improves considerably on Global EMOS. MEMOS

Table 1: Mean CRPS and mean AE for raw and postprocessed ensemble forecasts of surface temperature at individual stations.

	CRPS	AE
Raw ECMWF	2.50	2.81
Global EMOS	1.79	2.49
Local EMOS	1.42	1.97
MEMOS	1.40	1.97




Figure 2: Verification rank and PIT histograms for raw and postprocessed ensemble forecasts of surface temperature at individual stations.

Table 2: Mean energy score in the Northern coastline example.

	Independence	ECC
Raw ECMWF	6.32	6.37
Global EMOS	5.40	5.24
Local EMOS	4.79	4.74
MEMOS	4.81	4.66

borrows information from neighbouring stations via the Markovian dependence structure of the GMRF approximations for the bias parameters, and the Bayesian implementation takes account of forecast uncertainty, which leads to further improvement in the predictive performance when measured by the CRPS.

Figure 2 shows the rank histogram for the Raw ECMWF ensemble along with the PIT histograms for Global EMOS, Local EMOS, and MEMOS, respectively. Instead of plotting all possible $m + 1 = 51$ bins in the histograms, we use a slightly lower resolution and aggregate consecutive ranks into 17 bins. The same resolution is applied to construct the PIT histograms, for better comparison. The rank histogram indicates heavy underdispersion of the Raw ECMWF ensemble and suggests a pronounced need for postprocessing. The PIT histograms for Local EMOS and Global EMOS are much improved, even though they remain indicative of underdispersion. MEMOS shows the most uniform PIT histogram, well in line with the ranking in terms of the mean CRPS in Table 1.

We now apply the Diebold-Mariano test to the key comparison between MEMOS and Local EMOS. For the time series of the daily mean of the CRPS, the tail probability is smaller than 0.01. However, for the daily means of the AE, the tail probability is 0.42 which is in line with the overall mean AE for MEMOS being equal to that for Local EMOS, see Table 1.

3.4 Results at several stations simultaneously

We turn to a multivariate example with eleven stations along the North Sea and Baltic Sea coastlines, the locations of which are illustrated in Figure 1, to compare the predictive performance of the Raw ECMWF ensemble and the postprocessed Global EMOS, Local EMOS, and MEMOS forecasts under both ECC and Independence structures. As noted, the Raw ECMWF ensemble is invariant under ECC. The multivariate Raw ECMWF, Global EMOS, and Local EMOS ensembles are of size $m = 50$, while the MEMOS ensemble is of size $N = mn = 5,000$, both under ECC and Independence structures, as described in Section 2.4. Table 2 shows the mean energy score (7) and Figure 3 the multivariate rank histograms in this example. Clearly, ECC improves on the Independence approach, except in the case of the Raw ECMWF ensemble, where the Independence assumption compensates for the severe underdispersion, essentially replacing one evil by another. Global EMOS improves on the Raw ECMWF forecast, Local EMOS outperforms Global EMOS, and MEMOS improves on the Local EMOS benchmark, under the ECC approach. For each binary comparison, the tail probability under the Diebold-Mariano test of equal predictive performance is ≤ 0.01 , except for the comparison between Local EMOS Independence and MEMOS Independence, where it is 0.31. Similar patterns can be seen in the multivariate rank histograms, where Global EMOS ECC and MEMOS ECC show the most nearly uniform histograms.




Figure 3: Multivariate rank histograms in the Northern coastline example.

4 Discussion

The MEMOS approach provides a spatially adaptive, Bayesian implementation of the basic EMOS technique, where spatially varying bias coefficients are modeled as GMRFs. Inference is performed in a computationally efficient fashion within the SPDE-INLA framework of [Rue et al. \[2009\]](#) and [Lindgren et al. \[2011\]](#). Through the Markovian dependence structure of the GMRFs, MEMOS borrows information from neighbouring stations at the estimation stage, and the Bayesian implementation takes account of estimation uncertainty. In our case study of probabilistic temperature forecasts over Germany based on the ECMWF ensemble, MEMOS performs well against the Raw ECMWF ensemble, Global EMOS, and Local EMOS at observation sites. Furthermore, MEMOS allows for statistically postprocessed predictive distributions at any desired location.

We now describe potential future work. [Scheuerer and Büermann \[2014\]](#) and [Scheuerer and König \[2014\]](#) developed an EMOS version that incorporates information about the short-term local climatology, and their idea could be combined with the MEMOS approach. Future versions of MEMOS also might allow for spatially varying predictive variances, as opposed to the current implementation, which assumes a spatially constant predictive variance. In such an extension, the variance σ^2 can be parameterized as a linear function of the ensemble variance, in the sense that $\text{var}(\varepsilon) = c + dS^2$ in (1), where $S^2 = \frac{1}{m-1} \sum_{k=1}^m (f_k - \bar{f})^2$ is the ensemble variance [[Gneiting et al., 2005](#)]. The nonhomogeneous error term accounts for the ensemble spread-skill relationship as well as possible under- or overdispersion [[Whitaker and Lough, 1998](#)]. Furthermore, MEMOS can be tailored to ensembles with non-exchangeable members, or ensembles with groups of exchangeable members. As a caveat, while these are extensions, they may or may not prove beneficial in forecast mode.

To account for dependencies in forecasts at several locations simultaneously, we have combined MEMOS with ECC. As an alternative, it may be possible to develop a version of the spatial EMOS technique of [Feldmann et al. \[2015\]](#) that applies to MEMOS. More generally, methods that account for spatial and temporal as well as inter-variable dependencies are in high demand. Recent developments in these directions include both parametric and non-parametric copula approaches [[Möller et al., 2013](#), [Scheffzik et al., 2013](#), [Wilks, 2015](#), [Scheffzik, 2015](#)].

Our version of MEMOS is tailored to weather quantities for which the conditional predictive distributions can be assumed to be Gaussian. As the SPDE-INLA methodology is not restricted to normally distributed responses, it may be possible to develop variants for other weather quantities, utilizing the respective univariate postprocessing approaches, such as the EMOS techniques proposed for wind speed [[Thorarinsdottir and Gneiting, 2010](#)] and precipitation [[Scheuerer, 2014](#)], respectively.

As ensemble forecasts continue to improve over time, [Hemri et al. \[2014\]](#) analyze the evolution of the difference in skill between the Raw ECMWF and statistically postprocessed forecasts for a time period covering the years 2002 to 2014. Perhaps surprisingly, they find that the gap in skill remains almost constant over time. This suggests that improvements in numerical weather prediction models themselves, and improvements by statistical postprocessing, are complementary. In this light, we anticipate that statistical postprocessing will continue to yield substantial benefits in weather prediction for decades to come.

Acknowledgements

We thank Kira Feldmann, Roman Scheffzik, and Michael Scheuerer for helpful discussions, and Finn Lindgren for invaluable assistance regarding the understanding of the SPDE methodology

and the use of the R-INLA package. Furthermore, we are grateful to the European Centre for Medium-Range Weather Forecasts (ECMWF) and the German Weather Service (DWD) for providing forecast and observation data, respectively. Annette Möller and Alex Lenkoski gratefully acknowledge support by the German Research Foundation (DFG) within the programme “Spatio-/Temporal Graphical Models and Applications in Image Analysis” (GRK 1653).

References

- O. Barndorff-Nielsen, J. Kent, and M. Sørensen. Normal variance-mean mixtures and z distributions. *International Statistical Review*, 50:145–159, 1982.
- M. Blangiardo, M. Cameletti, G. Baio, and H. Rue. Spatial and spatio-temporal models with R-INLA. *Spatial and Spatio-Temporal Epidemiology*, 4:33–49, 2013.
- A. P. Dawid. Statistical theory: The prequential approach (with discussion). *Journal of the Royal Statistical Society Series A*, 147:278–292, 1984.
- A. F. Di Narzo and D. Cocchi. A Bayesian hierarchical approach to ensemble weather forecasting. *Journal of the Royal Statistical Society Series C*, 59:405–422, 2010.
- F. X. Diebold and R. S. Mariano. Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13:253–263, 1995.
- K. Feldmann, M. Scheuerer, and T. L. Thorarinsdottir. Spatial postprocessing of ensemble forecasts for temperature using nonhomogeneous Gaussian regression. *Monthly Weather Review*, 143:955–971, 2015.
- Y. Gel, A. E. Raftery, and T. Gneiting. Calibrated probabilistic mesoscale weather field forecasting: The geostatistical output perturbation (GOP) method (with discussion and rejoinder). *Journal of the American Statistical Association*, 99:575–590, 2004.
- B. Glahn, K. Gilbert, R. Cosgrove, D. P. Ruth, and K. Sheets. The gridding of MOS. *Weather and Forecasting*, 24:520–529, 2009.
- T. Gneiting. Making and evaluating point forecasts. *Journal of the American Statistical Association*, 106:746–762, 2011.
- T. Gneiting and M. Katzfuss. Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1:125–151, 2014.
- T. Gneiting and A. E. Raftery. Weather forecasting with ensemble methods. *Science*, 310:248–249, 2005.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378, 2007.
- T. Gneiting, A. Raftery, A. Westveld, and T. Goldman. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133:1098–1118, 2005.
- T. Gneiting, F. Balabdaoui, and A. E. Raftery. Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society Series B*, 69:243–268, 2007.

- T. Gneiting, L. I. Stanberry, E. P. Gritmit, L. Held, and N. A. Johnson. Assessing probabilistic forecasts of multivariate quantities, with applications to ensemble predictions of surface winds (with discussion and rejoinder). *Test*, 17:211–264, 2008.
- E. P. Gritmit, T. Gneiting, V. J. Berrocal, and N. A. Johnson. The continuous ranked probability score for circular variables and its application to mesoscale forecast ensemble verification. *Quarterly Journal of the Royal Meteorological Society*, 132:2925–2942, 2006.
- P. Guttorp and T. Gneiting. Studies in the history of probability and statistics XLIX: On the Matérn correlation family. *Biometrika*, 93:989–995, 2008.
- S. Hemri, M. Scheuerer, F. Pappenberger, K. Bogner, and T. Haiden. Trends in the predictive performance of raw ensemble weather forecasts. *Geophysical Research Letters*, 41:9197–9205, 2014.
- A. S. Hering and M. G. Genton. Comparing spatial predictions. *Technometrics*, 53:414–425, 2011.
- W. Kleiber, A. E. Raftery, J. Baars, T. Gneiting, C. Mass, and E. P. Gritmit. Locally calibrated probabilistic temperature forecasting using geostatistical model averaging and local Bayesian model averaging. *Monthly Weather Review*, 139:2630–2649, 2011a.
- W. Kleiber, A. E. Raftery, and T. Gneiting. Geostatistical model averaging for locally calibrated probabilistic quantitative precipitation forecasting. *Journal of the American Statistical Association*, 106:1291–1303, 2011b.
- M. Leutbecher and T. N. Palmer. Ensemble forecasting. *Journal of Computational Physics*, 227:3515–3539, 2008.
- F. Lindgren and H. Rue. Bayesian spatial modelling with R-INLA. *Journal of Statistical Software*, 63, 2015.
- F. Lindgren, H. Rue, and J. Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach (with discussion). *Journal of the Royal Statistical Society Series B*, 73:423–498, 2011.
- C. F. Mass, J. Baars, G. Wedam, E. Gritmit, and R. Steed. Removal of systematic model bias on a grid. *Weather and Forecasting*, 23:438–459, 2008.
- B. Matérn. *Spatial Variation*. Springer, 2nd edition, 1986.
- A. Möller, A. Lenkoski, and T. L. Thorarinsdottir. Multivariate probabilistic forecasting using Bayesian model averaging and copulas. *Quarterly Journal of the Royal Meteorological Society*, 139:982–991, 2013.
- F. Molteni, R. Buizza, T. N. Palmer, and T. Petroliagis. The new ECMWF ensemble prediction system: Methodology and validation. *Quarterly Journal of the Royal Meteorological Society*, 122:73–119, 1996.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013. URL <http://www.R-project.org/>.
- A. E. Raftery, T. Gneiting, F. Balabdaoui, and M. Polakowski. Using Bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, 133:1155–1174, 2005.

- H. Rue and L. Held. *Gaussian Markov Random Fields. Theory and Applications*. Chapman and Hall/CRC, 2005.
- H. Rue, S. Martino, and N. Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximation (with discussion). *Journal of the Royal Statistical Society Series B*, 71:319–392, 2009.
- R. Schefzik. Multivariate discrete copulas, with applications in probabilistic weather forecasting. *Publications de l’Institut de Statistique de l’Université de Paris*, 59:87–116, 2015.
- R. Schefzik, T. L. Thorarinsdottir, and T. Gneiting. Uncertainty quantification in complex simulation models using ensemble copula coupling. *Statistical Science*, 28:616–640, 2013.
- M. Scheuerer. Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Quarterly Journal of the Royal Meteorological Society*, 140:1086–1096, 2014.
- M. Scheuerer and L. Büermann. Spatially adaptive post-processing of ensemble forecasts for temperature. *Journal of the Royal Statistical Society Series C*, 63:405–422, 2014.
- M. Scheuerer and G. König. Gridded, locally calibrated, probabilistic temperature forecasts based on ensemble model output statistics. *Quarterly Journal of the Royal Meteorological Society*, 140:2582–2590, 2014.
- T. L. Thorarinsdottir and T. Gneiting. Probabilistic forecasts of wind speed: Ensemble model output statistics using heteroskedastic censored regression. *Journal of the Royal Statistical Society Series A*, 173:371–388, 2010.
- J. A. Vrugt, C. G. H. Diks, and M. P. Clark. Ensemble Bayesian model averaging using Markov chain Monte Carlo sampling. *Environmental Fluid Mechanics*, 8:579–595, 2008.
- J. S. Whitaker and A. F. Lough. The relationship between ensemble spread and ensemble mean skill. *Monthly Weather Review*, 126:3292–3302, 1998.
- D. S. Wilks. *Statistical Methods in the Atmospheric Sciences*. Elsevier Academic Press, 3rd edition, 2011.
- D. S. Wilks. Multivariate ensemble model output statistics using empirical copulas. *Quarterly Journal of the Royal Meteorological Society*, 141:945–952, 2015.
- D. S. Wilks and T. M. Hamill. Comparison of ensemble-MOS methods using GFS reforecasts. *Monthly Weather Review*, 135:2379–2390, 2007.